

ECS Configuration Change Request

Page 1 of

Page(s)

1. Originator Byron V. Peters	2. Log Date: 17 Jul 00	3. CCR #: 00-0739	4. Rev: -	5. Tel: 301/883-4077	6. Rm #: 2101C	7. Dept. SED
8. CCR Title: Release COTS test executable for IRIX 6.2 Challenges with Gigabit Ethernet						
9. Originator Signature/Date <i>Byron V. Peters</i> 7/17/00		10. Class II	11. Type: CCR	12. Need Date: 7/17/00		
13. Office Manager Signature/Date <i>K.P. Arguebar</i> 07/17/00		14. Category of Change: Update ECS Baseline Doc.		15. Priority: (If "Emergency" fill in Block 28). Emergency		
16. Documentation/Drawings Impacted: 010 -TDA-004 Rev 7 (will be updated on next SGI patch release) 9/11		17. Schedule Impact:		18. CI(s) Affected: ICL, DADS		
19. Release Affected by this Change: 5A, 5B		20. Date due to Customer:		21. Estimated Cost: None - Under 100K		
22. Source Reference: <input type="checkbox"/> NCR (attach) <input type="checkbox"/> Action Item <input type="checkbox"/> Tech Ref. <input type="checkbox"/> GSFC <input type="checkbox"/> Other: SMC 000 00000 2539						
23. Problem: (use additional sheets if necessary) In the process of installing the Cisco Catalyst 6500 Ethernet switch at NSIDC, severe errors on the interface input, throughput problems and system crashes were encountered on SGI systems running with the Gigabit Ethernet vfe interfaces under IRIX 6.2. This does not seem to be a problem with IRIX 6.5.						
24. Proposed Solution: (use additional sheets if necessary) SGI has a patch patchSGI00003563 that replaces the vfe binary. The patch has been field tested at NSIDC and appears to fix the it needs to be installed wherever there are SGI Challenges with vfe Gigabit Ethernet interfaces running IRIX 6.2. In order to expedite the process, it is requested to release this OS patch as a test executable which will be rolled up into the next patch release.						
25. Alternate Solution: (use additional sheets if necessary) Wait till the patch has been fully regression tested at Landover and then release.						
26. Consequences if Change(s) are not approved: (use additional sheets if necessary) NSIDC operations will continue to be down.						
27. Justification for Emergency (If Block 15 is "Emergency"): NSIDC is out of service until the patch is released.						
28. Site(s) Affected: <input type="checkbox"/> EDF <input type="checkbox"/> PVC <input type="checkbox"/> VATC <input type="checkbox"/> EDC <input checked="" type="checkbox"/> GSFC <input checked="" type="checkbox"/> LaRC <input checked="" type="checkbox"/> NSIDC <input checked="" type="checkbox"/> SMC <input type="checkbox"/> AK <input type="checkbox"/> JPL <input type="checkbox"/> EOC <input type="checkbox"/> IDG Test Cell <input type="checkbox"/> Other						
29. Board Comments:			30. Work Assigned To:		31. CCR Closed Date:	
32. EDF/SCDV CCB Chair (Sign/Date):		Disposition: Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ECS				
33. M&O CCB Chair (Sign/Date): <i>[Signature]</i> 7/18/00		Disposition: Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ECS				
34. ECS CCB Chair (Sign/Date):		Disposition: Approved App/Com. Disapproved Withdraw Fwd/ESDIS ERB Fwd/ESDIS				

CM01JA00

ECS/EDF/SCDV/M&O

ORIGINAL

00-0739

[HOME](#) [SEARCH](#) [HOW TO BUY](#) [products](#) [solutions](#) [developers](#) [services & support](#) [serious fun](#)Hello, **bpeters!**

IRIX

Supportfolio
online[News](#) [MyProfile](#) [Feedback](#) [Search](#) [SiteMap](#) [Help](#)**Patch**

IRIX Patches

GO

[Patch Sets](#) | [Maintenance Release](#) | [Find Patches](#) | [Patch News](#) | [Security Headquarters](#)

Patch 3563 [IRIX 6.2] - 100BT VFE driver roll-up #5 for IRIX 6.2

INDEX

- [Relations](#)
- [Release Notes](#)
- [Inst Subsystem Requirements](#)
- [Inst Subsystem Checksums](#)
- [Inst Subsystem File Listings](#)
- [Download Patch](#)

RELATIONS

This patch replaces the following patches:

1623, 1867, 2182, 2441, 2948, 3495

This patch has no known incompatibilities with other patches.

This patch fixes the following bugs:

426883 - No more than 6 ipaliases possible on 100baseT vfe interfaces
454773 - BUG Poor NFS performance over VME based 100BaseT Boards
467850 - vfe0 drops, doesn't come back
499959 - vfe interrupt hold too long
516376 - strange panic in lf_free through vfe code
564518 - vfe panic in mutex_spinlock()
571220 - vfe interface is slow responding to initial ping and arp request
651621 - vfe DMA unfair for 2nd port on each board

RELEASE NOTES

1. Patch SG0003563 Release Note

This release note describes patch SG0003563 to IRIX 6.2.

Make sure that the Network Patch SG0002673 or the latest Network Rollover Patch has installed. Also make sure that

00-0739

the OS patch SG0003252 or the latest OS Rollover Patch has installed BEFORE installing this patch SG0003563.

Patch SG0003563 replaces patches(es) SG0003495, SG0002948, SG0002441, SG0001623, SG0001867, SG0002182.

1.1 Supported Hardware Platforms

This patch contains bug fixes for all IP19, IP21, IP25 systems. The software cannot be installed on other configurations.

1.2 Supported Software Platforms

This patch contains bug fixes for VFE 2.1 (version 1233254020) on a system running IRIX 6.2. The software cannot be installed on other configurations.

1.3 Bugs Fixed by Patch SG0003563

This patch contains fixes for the following bugs in IRIX 6.2. Bug numbers from Silicon Graphics bug tracking system are included for reference.

- o Bug 467850, 516376 - VFE interface accumulated high number of ierrs, causing the route to drop, or system hang in some cases. The problem was the VFE driver incorrectly read in the network statistic registers from the board, causing 'netstat' program to report accumulated ierrors from vfe interface. Not reading the registers correctly also caused them not being cleared. Eventually the high number of ierrs caused routing problem, or slow network performance. Registers not being cleared also generated non-stop STATS_OVERFLOW interrupt to happen, which in turn caused the system to hang. Fix in this patch addressed this problem.
- o Bug 499959 - VFE interrupts hold too long, affecting Real Time Applications. Using tools like WindView, it's noticed with vfe interface, several sockd spikes would occur, with silence time average to ~50ms. Frequency of such was found related to when the watchdog timer routine in vfe driver was called.

Problem was due to in this routine it executed MII reads more than necessary, and not efficiently enough, hence causing kernel not preemptable for too long. Fixing the algorithm reduced the time down to ~7ms.

- o Bug 426883 - No more than 6 ipaliases possible on 100BaseT vfe interfaces. The problem was originally related to HA FailSafe script, where it couldn't handle a case where the vfe driver would get an error message like "vfe: Adapter Chk Interrupt.. = Ack Count Error". Under such case, the driver would automatically recover from it. However the FailSafe script wasn't equipped to deal with error conditions, and therefore timed out.

00-0739

The script has since then been modified to deal with this condition. Later the source of the problem has been found, which makes of this patch for. Basically the adapter check message was related to the chip complaining inconsistency of the number of interrupts ack'ed by the driver. That is, the driver ack'ed more interrupts than the chip thinks it has pending. The was due to a race condition in driver during interrupt handling, in clearing up a flag, hence caused the problem.

- o Bug 454773 -Poor NFS performance. Poor NFS performance was both noted with NFS v.2 and NFS v.3. Large numbers of retransmit and time-out were also noted. One of the problems was found in the driver, that the interrupt holdoff timer was set to be too short, causing the interrupts happened quicker then it could handle. In the case where packets coming in succession, it stood the chance of missing interrupts. This is especially obvious when testing with NFS v.2.

Another problem was found relating to NFS v.3, due to its features of 32K read/write transfers, and the readahead, which in turn causes several 32K read transfers in rapid succession. However the driver set the transmit queue number to be just 32, so one 32K read would fill the output queue enough. Increased this number addressed the problem.

- o This patch also contains a fix for a problem, that the timeout for checking the link status (PHY output) was set to be too long, causing the system hanging longer than necessary, before the timer expired. The system might look like hanging while in fact it's waiting for the timer to expire. This patch has the fix to reduce the value of the said timer.
- o Bug 571220 - vfe interface is slow responding to initial ping and arp request. the fix is on the vfe interrupt handler for the ARP msg. if the very previous msg is ARP request, it gets the msg from queue and prepares the DMA list to send the ARP respond right away.

our test result shows 2-3% performance increase besides resolved the orig problem.

- o Bug 564518 - vfe panic in mutex_spinlock(). if_vfeintr() has been modified to reclaim the transmit mbufs ONLY after the interrupt handler released the IFNET_spin_lock, and the uv_dma_done() is called ONLY when the a DMA is done, it should not been called by 5 other error interrupts. This fix is compiled and tested with the latest Network patch2673 and OS patch3252. Make sure that these 2 patches or the latest Network, OS rollover patches have installed before installing this patch.
- o Bug 651621 - vfe DMA unfair for 2nd port on each board. The fix is to equally distribute dma list for transmit and receive, when the 2nd port is detected up.

1.4 Subsystems Included in Patch SG0003563

00-0739

This patch release includes these subsystems:

- o patchSG0003563.vfe_sw
- o patchSG0003563.idb

1.5 Installation Instructions

Because you want to install only the patches for problems you have encountered, patch software is not installed by default. After reading the descriptions of the bugs fixed in this patch (see Section 1.3), determine the patches that meet your specific needs.

If, after reading Sections 1.1 and 1.2 of these release notes, you are unsure whether your hardware and software meet the requirements for installing a particular patch, run `inst`. The `inst` program does not allow you to install patches that are incompatible with your hardware or software.

Patch software is installed like any other Silicon Graphics software product. Follow the instructions in your Software

Installation Administrator's Guide to bring up the miniroot form of the software installation tools.

Follow these steps to select a patch for installation:

1. At the `Inst>` prompt, type

`install patchSGxxxxxxx`

where `xxxxxxx` is the patch number.
2. Initiate the installation sequence. Type

`Inst> go`
3. You may find that two patches have been marked as incompatible. (The installation tools reject an installation request if an incompatibility is detected.) If this occurs, you must deselect one of the patches.

`Inst> keep patchSGxxxxxxx`

where `xxxxxxx` is the patch number.
4. After completing the installation process, exit the `inst` program by typing

`Inst> quit`

1.6 Patch Removal Instructions

To remove a patch, use the versions `remove` command as you

00-0739

would for any other software subsystem. The removal process reinstates the original version of software unless you have specifically removed the patch history from your system.

versions remove patchSGxxxxxxx

where xxxxxxx is the patch number.

To keep a patch but increase your disk space, use the versions removehist command to remove the patch history.

versions removehist patchSGxxxxxxx

where xxxxxxx is the patch number.

1.7 Known Problems

INST SUBSYSTEM REQUIREMENTS

No Requirements Information Available.

INST SUBSYSTEM CHECKSUMS

These checksums help to provide a 'signature' for the patch inst image which can be used to authenticate other inst images. You can obtain this kind of output by running **sum -r** on the image (from the command line):

22075	3	patchSG0003563.idb
48025	17	patch/README.patch.3563
12613	2	patchSG0003563

INST SUBSYSTEM FILE LISTINGS

The following lists the files which get installed from each subsystem in the patch:

patchSG0003563.vfe_sw.eoe

- usr/cpu/sysgen/IP19boot/if_vfe.a
- usr/cpu/sysgen/IP21boot/if_vfe.a
- usr/cpu/sysgen/IP25boot/if_vfe.a
- usr/relnotes/patchSG0003563/TC
- usr/relnotes/patchSG0003563/ch1.z

DOWNLOAD PATCH

The size of the patch is 133Kb. [Download the patch.](#)

If you have support contract, you may log a call electronically with [Supportfolio Connect](#) to get this patch shipped to you.

00-0739

Maintained by supportlib@sgi.com

This site is maintained by [Electronic Services](#), Global Customer Service, Silicon Graphics, Inc.
All materials here are © [Silicon Graphics, Inc.](#) 1997-2000. All rights reserved.
All other trademarks are the property of their [respective owners](#).